

US Patent 5,611,049, Claim 16:

16. A network of digital computers that includes a first plurality of client sites which request access to a stored dataset that is stored at a location that can be accessed through the network, the network comprising:

**a second plurality of NDC sites**, the stored dataset whose access is requested by the client sites being stored at **an NDC server terminator site**, a request from the client sites for access to the stored dataset being received by **a third plurality of NDC client terminator sites**, each NDC site including:

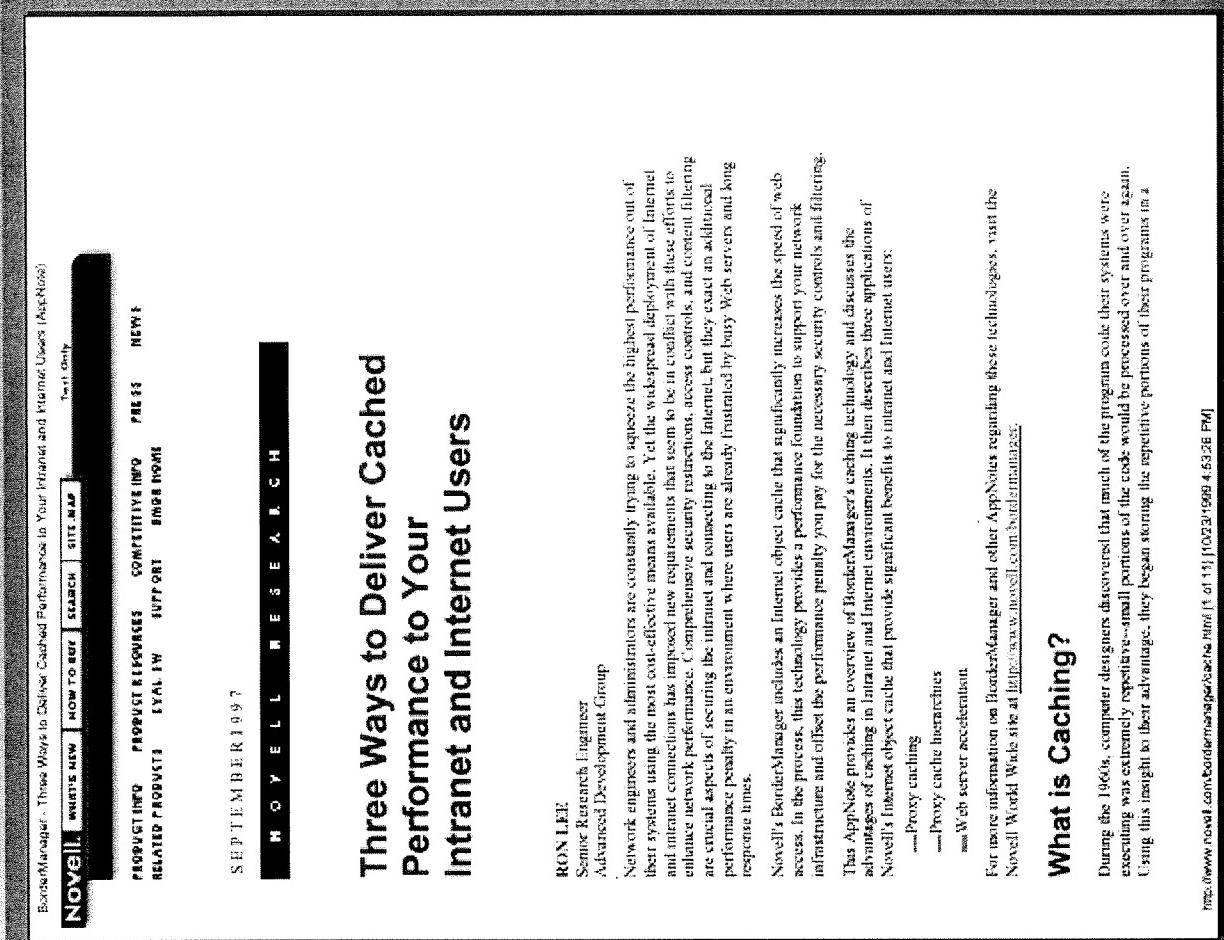
- (a) an NDC that has an NDC buffer;
  - (b) means for the NDC to receive the request to access the stored dataset;
  - (c) means for the NDC to check the NDC buffer at this NDC site to determine if a projected image of data requested from the stored dataset is already present there, wherein:

- i. if the NDC buffer of this NDC site does not contain a projected image of all data requested from the stored dataset, and if this NDC site is not the NDC server terminator site for the stored dataset, the NDC includes means for transmitting a request for data from this NDC site downstream to another NDC site closer to the NDC server terminator site for the stored dataset than the present NDC site.

- ii. if the NDC buffer of this NDC site does not contain a projected image of all data requested from the stored dataset, and if this NDC site is the NDC server terminator site for the stored dataset, the NDC including means for accessing the stored dataset to project an image of the requested data into the buffer of this NDC; and

- iii. if the NDC buffer of an NDC site contains a projected image of all requested data, the NDC including means for returning the data requested from this NDC site upstream to the NDC site from which this NDC site received the request, whereby through a succession of such returns of data from one NDC site to the next upstream NDC site the requested data ultimately arrives at the NDC client terminator site, each NDC site that returns data upstream to the requesting NDC site retaining a copy of the returned data that the returning NDC site may subsequently transmit to an NDC site other than the NDC site to which the returning NDC site first returned the data, whereby images of the stored dataset may be projected concurrently from a single NDC

(d) means for the NDC client terminator site to return the requested data to the client site that requested access to the stored dataset.



NCT010702

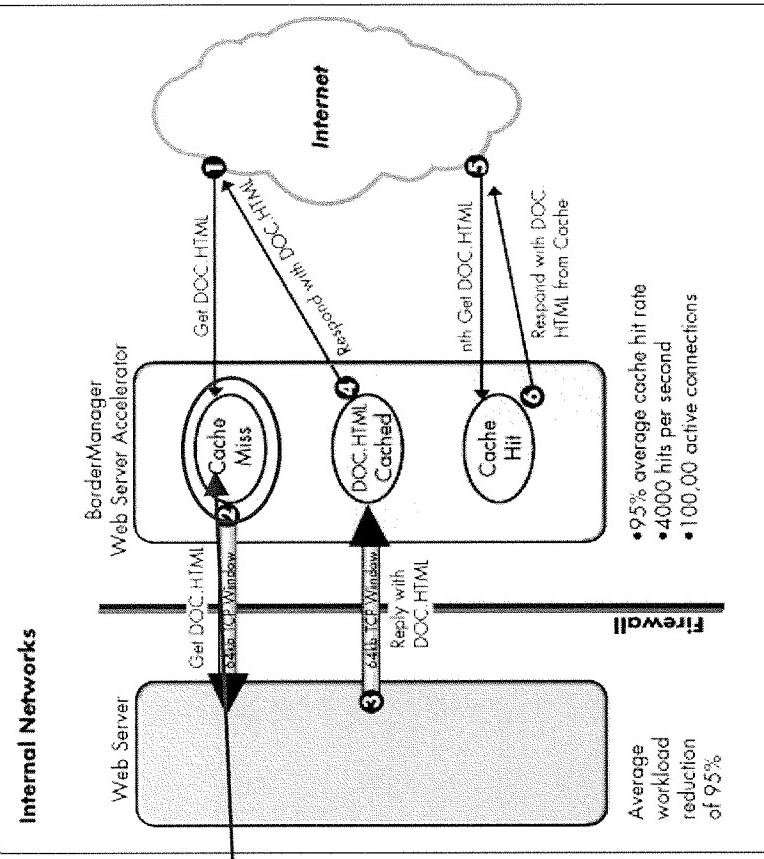
49.16

## Web Server Acceleration

Web servers can be a bottleneck in your intranet or Internet infrastructures. Typical web servers quickly run out of connection capacity and tend to produce slow response times. In sites where performance is important, the only options usually considered are to upgrade to a more expensive web server system or to split the content set across multiple web servers. Neither of these options make sense when caching offers such an elegant, cost-effective means to overcome the problem.

Configured as a web server accelerator, Novell's Internet object cache eliminates the web server bottleneck by placing a dedicated cache in front of the web server and handling requests for all of the web server's cachable content directly from its own cache. Caching is the obvious solution because typical web sites are constructed with approximately 95-100 percent cachable content. Once this material is fetched from the web server and cached in the web server accelerator, the accelerator can handle all of the requests for that content. This leaves the small percentage of dynamic requests to be "passed through" the accelerator for the origin web server to process (see Figure 8).

Figure 8: The web server accelerator offloads over 90 percent of the web server's workload and responds to requests at cached speeds.



(a) an NDC that has an NDC buffer;

(b) means for the NDC to receive the request to access the stored dataset;

(c) means for the NDC to check the NDC buffer at this NDC site to determine if a projected image of data requested from the stored dataset is already present there, wherein:

- if the NDC buffer of this NDC site does not contain a projected image of all data requested from the stored dataset, and if this NDC site is not the NDC server terminator site for the stored dataset, the NDC includes means for transmitting a request for data from this NDC site downstream to another NDC site closer to the NDC server terminator site for the stored dataset than the present NDC site;
- if the NDC buffer of this NDC site does not contain a projected image of all data requested from the stored dataset, and if this NDC site is the NDC server terminator site for the stored dataset, the NDC including means for accessing the stored dataset to project an image of the requested data into the buffer of this NDC; and

1. A browser issues a request for a file named 'DOC.HTML'. This request is received by the web server accelerator. In this case, the request results in a "cache miss" because the web server accelerator has never serviced a request for that document before.

2. The web server accelerator initiates a request for 'DOC.HTML' from your web server on behalf of the browser.

3. The origin web server responds to the web server accelerator's request by sending 'DOC.HTML'. This transmission is much faster than a response to a browser due to the web server accelerator's optimized receive bandwidth that can receive up to 64KB at one time and that stays open to receive multiple responses. The web server accelerator responds to the original browser request with 'DOC.HTML'.

4. Now when the same browser (or any other browser) issues a request for 'DOC.HTML', the request results in a "cache hit" because the web server accelerator has kept a copy of the document in its cache.

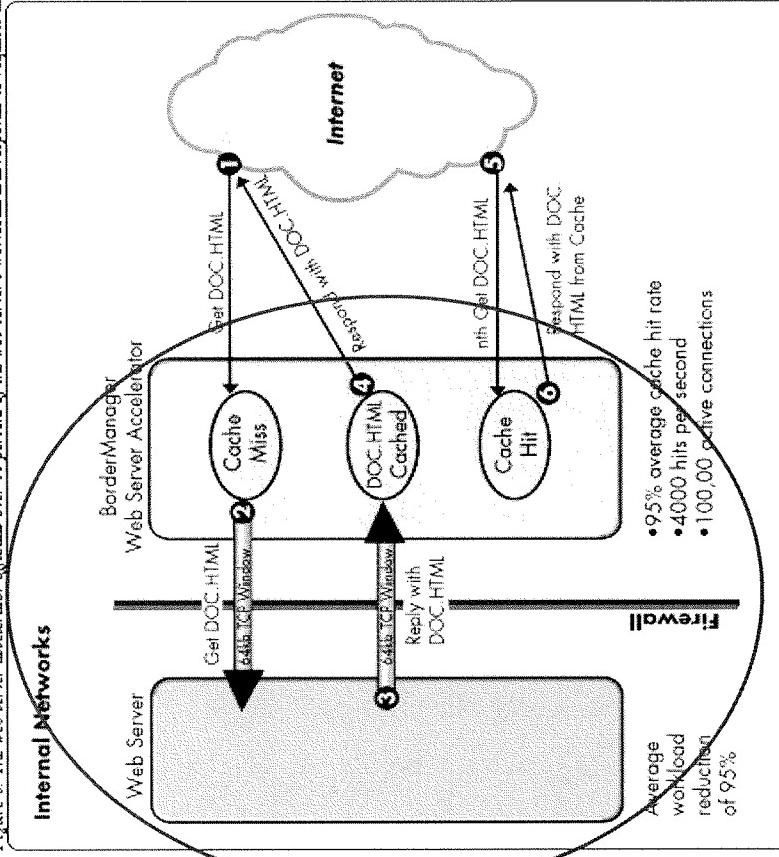
5. In this case, the web server accelerator replies immediately to the browser request because it has 'DOC.HTML' in cache. The proxy's response eliminates the need to fetch the document again from the origin web server.

## Web Server Acceleration

Web servers can be a bottleneck in your intranet or Internet infrastructures. Typical web servers quickly run out of connection capacity and tend to produce slow response times. In sites where performance is important, the only options usually considered are to upgrade to a more expensive web server system or to split the content set across multiple web servers. Neither of these options make sense when caching offers such an elegant, cost-effective means to overcome the problem.

Configured as a web server accelerator, Novell's Internet object cache eliminates the web server bottleneck by placing a dedicated cache in front of the web server and handling requests for all of the web server's cacheable content directly from its own cache. Caching is the obvious solution because typical web sites are constructed with approximately 95-100 percent cacheable content. Once this material is fetched from the web server and cached in the web server accelerator, the accelerator can handle all of the requests for that content. This leaves the small percentage of dynamic requests to be "passed through" the accelerator for the origin web server to process (see Figure 8).

**Figure 8: The web server accelerator offloads over 90 percent of the web server's workload and responds to requests at cached speed.**



(a) an NDC that has an NDC buffer;

(b) means for the NDC to receive the request to access the stored data;

(c) means for the NDC to check the NDC buffer at this NDC site to determine if a projected image of data requested from the stored dataset is already present there, wherein:

- i. if the NDC buffer of this NDC site does not contain a projected image of all data requested from the stored dataset, and if this NDC site is not the NDC server terminator site for the stored dataset, the NDC includes the NDC server terminator site for the stored dataset from this NDC site means for transmitting a request for data from this NDC site downstream to another NDC site closer to the NDC server terminator site for the stored dataset than the present NDC site;

- ii. if the NDC buffer of this NDC site does not contain a projected image of all data requested from the stored dataset, and if this NDC site is the NDC server terminator site for the stored dataset, the NDC including means for accessing the stored dataset to project an image of the requested data into the buffer of this NDC; and

1. A browser issues a request for a file named DOC.HTML. This request is received by the web server accelerator. In this case, the request results in a "cache miss" because the web server accelerator has never serviced a request for that document before.

2. The web server accelerator initiates a request for DOC.HTML from your web server on behalf of the browser.

3. The origin web server responds to the web server accelerator's request by sending DOC.HTML. This transmission is much faster than a response to a browser due to the web server accelerator's optimized receive window that can receive up to 64KB at one time and that stays open to receive multiple responses. The web server accelerator responds to the original browser request with DOC.HTML.

4. The web server accelerator replies immediately to the browser request because it has DOC.HTML in its cache.

5. Now when the same browser (or any other browser) issues a request for DOC.HTML, the request results in a "Cache Hit" because the web server accelerator has kept a copy of the document in its cache.

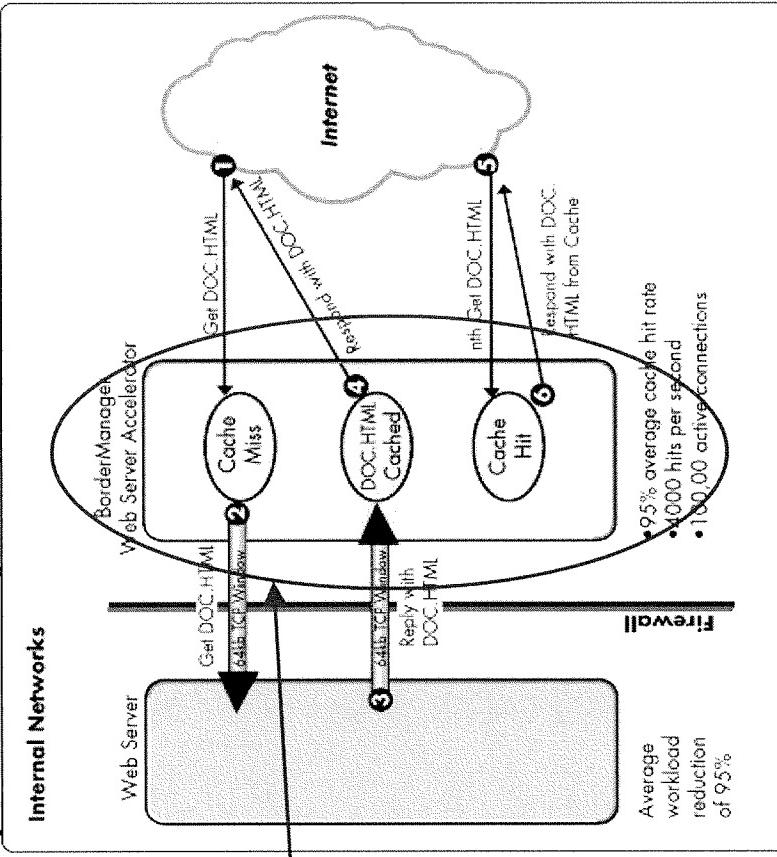
6. In this case, the web server accelerator replies immediately to the browser request because it has DOC.HTML in cache. The proxy's response eliminates the need to fetch the document again from the origin web server.

## Web Server Acceleration

Web servers can be a bottleneck in your intranet or Internet infrastructures. Typical web servers quickly run out of connection capacity and tend to produce slow response times. In sites where performance is important, the only options usually considered are to upgrade to a more expensive web server system or to split the content set across multiple web servers. Neither of these options make sense when caching offers such an elegant, cost-effective means to overcome the problem.

Configured as a web server accelerator, Novell's Internet object cache eliminates the web server bottleneck by placing a dedicated cache in front of the web server, archiving requests for all of the web server's cacheable content directly from its own cache. Caching is the obvious solution because typical web sites are constructed with approximately 95-100 percent cacheable content. Once this material is fetched from the web server and cached in the web server accelerator, the accelerator can handle all of the requests for that content. This leaves the small percentage of dynamic requests to be "passed through" the accelerator for the origin web server to process (see Figure 8).

**Figure 8: The web server accelerator offloads over 90 percent of the web server's workload and responds to requests at cache speed.**



(a) an NDC that has an NDC buffer;

- (b) means for the NDC to receive the request to access the stored data;
- (c) means for the NDC to check the NDC buffer at this NDC site to determine if a projected image of data requested from the stored dataset is already present there, wherein:

i. if the NDC buffer of this NDC site does not contain a projected image of all data requested from the stored dataset, and if this NDC site is not the NDC server terminator site for the stored dataset, the NDC includes the NDC server terminator site for the stored dataset, the NDC means for transmitting a request for data from this NDC site downstream to another NDC site closer to the NDC server terminator site for the stored dataset than the present NDC site;

ii. if the NDC buffer of this NDC site does not contain a projected image of all data requested from the stored dataset, and if this NDC site is the NDC server terminator site for the stored dataset, the NDC including means for accessing the stored dataset to project an image of the requested data into the buffer of this NDC; and

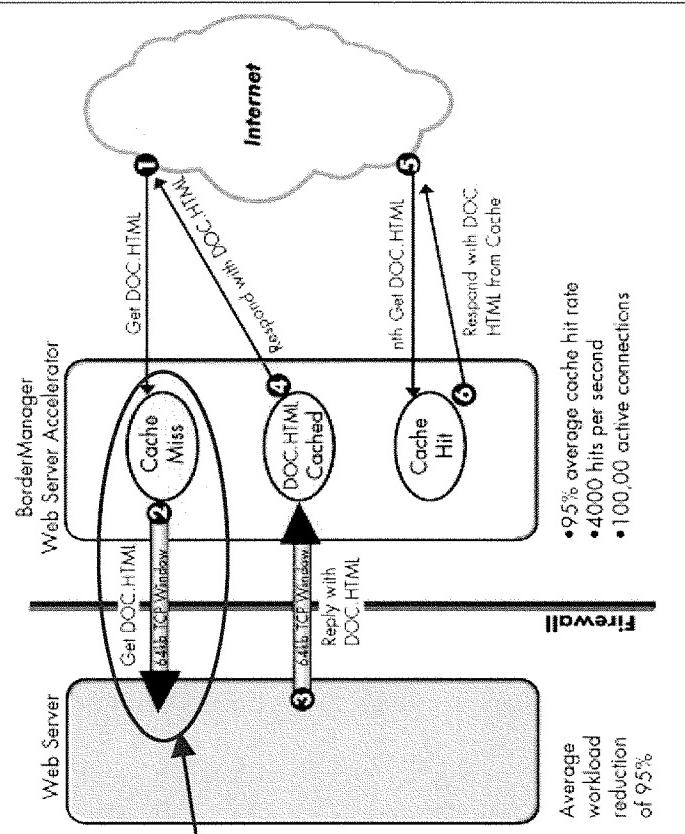
1. A browser issues a request for a file named DOC.HTML. This request is received by the web server accelerator. In this case, the request results in a "cache miss" because the web server accelerator has never serviced a request for that document before.
2. The web server accelerator initiates a request for DOC.HTML from your web server on behalf of the browser.
3. The origin web server responds to the web server accelerator's request by sending DOC.HTML. This transmission is much faster than a response to a browser due to the web server accelerator's optimized, recursive vendor that can receive up to 64KB at one time and that stays open to receive multiple responses. The web server accelerator responds to the original browser request with DOC.HTML.
4. The web server accelerator responds to the original browser request with DOC.HTML.
5. Now when the same browser (or any other browser) issues a request for DOC.HTML, the request results in a "cache hit" because the web server accelerator has kept a copy of the document in its cache.
6. In this case, the web server accelerator replies immediately to the browser request because it has DOC.HTML in cache. The proxy's response eliminates the need to fetch the document again from the origin web server.

## Web Server Acceleration

Web servers can be a bottleneck in your intranet or Internet infrastructures. Typical web servers quickly run out of connection capacity and tend to produce slow response times. In sites where performance is important, the only options usually considered are to upgrade to a more expensive web server system or to split the content set across multiple web servers. Neither of these options make sense when caching offers such an elegant, cost-effective means to overcome the problem.

Configured as a web server accelerator, Novell's Internet object cache eliminates the web server bottleneck by placing a dedicated cache in front of the web server and handling requests for all of the web server's cacheable content directly from its own cache. Caching is the obvious solution because typical web sites are constructed with approximately 95-100 percent cacheable content. Once this material is fetched from the web server and cached in the web server accelerator, the accelerator can handle all of the requests for that content. This leaves the small percentage of dynamic requests to be "passed through" the accelerator for the origin web server to process (see Figure 8).

**Figure 8: The web server accelerator offloads over 90 percent of the web server's workload and responds to requests at cache speeds.**



1. A browser issues a request for a file named DOC.HTML. This request is received by the web server accelerator. In this case, the request results in a "cache miss" because the web server accelerator has never serviced a request for that document before.

2. The web server accelerator initiates a request for DOC.HTML from your web server on behalf of the browser.

3. The origin web server responds to the web server accelerator's request by sending DOC.HTML. This transmission is much faster than a response to a browser due to the web server accelerator's optimized receive window that can receive up to 64KB at one time and that stays open to receive multiple responses. The web server accelerator responds to the original browser request with DOC.HTML.

4. Now when the same browser (or any other browser) issues a request for DOC.HTML, the request results in a "cache hit" because the web server accelerator has kept a copy of the document in its cache.

5. In this case, the web server accelerator replies immediately to the browser request because it has DOC.HTML in cache. The proxy's response eliminates the need to fetch the document again from the origin web server.

(a) an NDC that has an NDC buffer;

(b) means for the NDC to receive the request to access the stored data;

(c) means for the NDC to check the NDC buffer at this NDC site to determine if a projected image of data requested from the stored dataset is already present there, wherein:

i. if the NDC buffer of this NDC site does not contain a projected image of all data requested from the stored dataset, and if this NDC site is not the NDC server terminator site for the stored dataset, the NDC includes means for transmitting a request for data from this NDC site downstream to another NDC site closer to the NDC server terminator site for the stored dataset than the present NDC site;

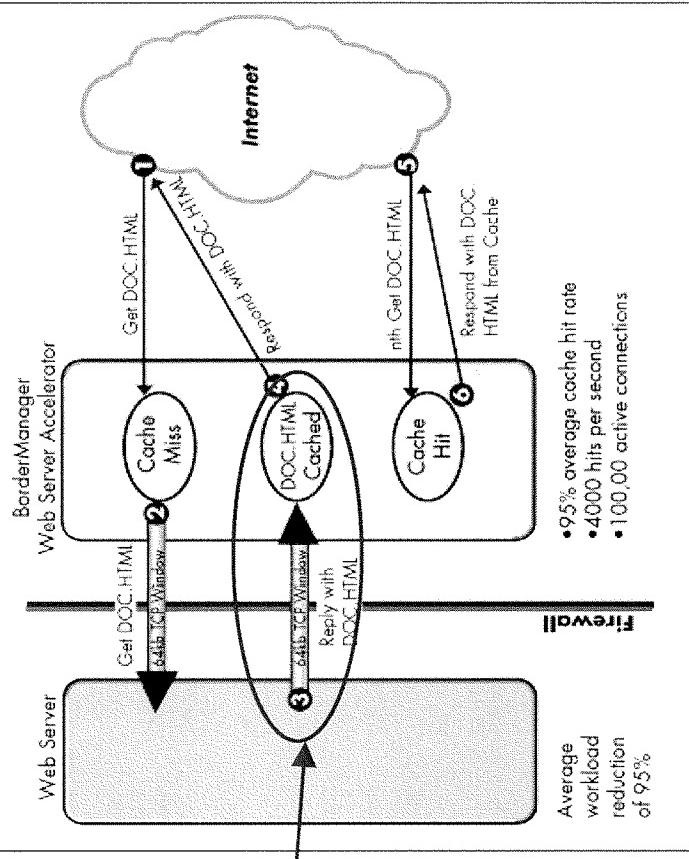
ii. if the NDC buffer of this NDC site does not contain a projected image of all data requested from the stored dataset, and if this NDC site is the NDC server terminator site for the stored dataset, the NDC including means for accessing the stored dataset to project an image of the requested data into the buffer of this NDC; and

## Web Server Acceleration

Web servers can be a bottleneck in your intranet or Internet infrastructures. Typical web servers quickly run out of connection capacity and tend to produce slow response times. In this where performance is important, the only options usually considered are to upgrade to a more expensive web server system or to split the content set across multiple web servers. Neither of these options make sense when caching offers such an elegant, cost-effective means to overcome the problem.

Configured as a web server accelerator, Novell's Internet object cache eliminates the web server bottleneck by placing a dedicated cache in front of the web server and handling requests for all of the web server's cacheable content directly from its own cache. Caching is the obvious solution because typical web sites are constructed with approximately 95-100 percent cacheable content. Once this material is fetched from the web server and cached in the web server accelerator, the accelerator can handle all of the requests for that content. This leaves the small percentage of dynamic requests to be "passed through" the accelerator for the origin web server to process (see Figure 8).

Figure 8. The web server accelerator offloads over 90 percent of the web server's workload and responds to requests at cache speed.



1. A browser issues a request for a file named DOC.HTML. This request is received by the web server accelerator. In this case, the request results in a "cache miss" because the web server accelerator has never serviced a request for that document before.

The web server accelerator initiates a request for DOC.HTML from your web server on behalf of the browser.

2. The origin web server responds to the web server accelerator's request by sending up DOC.HTML. This transmission is much faster than a response to a browser due to the web server accelerator's optimized receive window that can receive up to 64KB at one time and that stays open to receive multiple responses. The web server accelerator then places DOC.HTML in its cache.

3. Now when the same browser (or any other browser) issues a request for DOC.HTML, the request results in a "cache hit" because the web server accelerator has kept a copy of the document in its cache.

4. The web server accelerator responds to the original browser request with DOC.HTML.
5. Now when the same browser (or any other browser) issues a request for DOC.HTML, the request results in a "cache hit" because the web server accelerator has kept a copy of the document in its cache.
6. In this case, the web server accelerator replies immediately to the browser request because it has DOC.HTML in cache. The proxy's response eliminates the need to fetch the document again from the origin web server.

(a) an NDC that has an NDC buffer;

(b) means for the NDC to receive the request to access the stored dataset;

(c) means for the NDC to check the NDC buffer at this NDC site to determine if a projected image of data requested from the stored dataset is already present there, wherein:

- i. if the NDC buffer of this NDC site does not contain a projected image of all data requested from the stored dataset, and if this NDC site is not the NDC server terminator site for the stored dataset, the NDC includes means for transmitting a request for data from this NDC site downstream to another NDC site closer to the NDC server terminator site for the stored dataset than the present NDC site;
- ii. if the NDC buffer of this NDC site does not contain a projected image of all data requested from the stored dataset, and if this NDC site is the NDC server terminator site for the stored dataset, the NDC including means for accessing the stored dataset to project an image of the requested data into the buffer of this NDC; and

iii. if the NDC buffer of this NDC site does contain a projected image of all data requested from the stored dataset, the NDC includes

means for transmitting a request for data from this NDC site downstream to another NDC site closer to the NDC server terminator site for the stored dataset than the present NDC site;